Machine learning characterisation of Alfvénic and sub-Alfvénic chirping and correlation with fast ion loss at NSTX

B. J. Q. Woods,[1, 2] V. N. Duarte,[3] E. D. Fredrickson,[3] N. N. Gorelenkov,[3] M. Podestà,[3] and R. G. L. Vann[2]

[1] *School of Mathematics, University of Leeds, Woodhouse Lane, Leeds, LS2 9JT, United Kingdom*

[2] *Department of Physics, York Plasma Institute, University of York, Heslington, York, YO10 5DD, United Kingdom*

[3] *Princeton Plasma Physics Laboratory, Princeton University, Princeton, NJ 08543, United States of America*

(Dated: 13 December 2019)

Abrupt large events in the Alfvénic and sub-Alfvénic frequency bands in tokamaks are typically correlated with increased fast ion loss. Here, machine learning is used to speed up the laborious process of characterizing the behaviour of magnetic perturbations from corresponding frequency spectrograms that are typically identified by humans. Analysis allows for comparison between different mode character (such as quiescent, fixed-frequency, chirping, avalanching) and plasma parameters obtained from the TRANSP code such as the ratio of the neutral beam injection (NBI) velocity and the Alfvén velocity ($v_{\mathrm{inj.}}/v_{\mathrm{A}}$), the q-profile, and the ratio of the neutral beam beta and the total plasma beta ($\beta_{\mathrm{beam},i}/\beta$). In agreement with previous work by Fredrickson *et al.* [Nucl. Fusion 2014, 54 093007], we find correlation between $\beta_{\mathrm{beam},i}$ and mode character. In addition, previously unknown correlations are found between moments of the spectrograms and mode character. Character transition from quiescent to non-quiescent behaviour for magnetic fluctuations in the 50 - 200 kHz frequency band is observed along the boundary $v_{\varphi} \lessgtr \frac{1}{4}(v_{\mathrm{inj.}} - 3v_{\mathrm{A}})$ where v_{φ} is the rotation velocity.

I. INTRODUCTION

Fast ion instabilities correlated with increased Alfvénic activity could prove to be a serious limitation to the nominal ITER performance;[1] these wave-particle instabilities transfer free energy between the particle distribution function and plasma waves in the system, in certain cases leading to sudden degradation of plasma performance and energy confinement. Abrupt large events (ALEs or mode avalanches) are characterised by magnetic perturbations in the plasma undergoing very rapid frequency change ('chirping') across a broadband of frequencies, and are directly correlated with large energetic particle (EP) losses[2,3]; understanding the parametric dependencies on these losses is essential for good plasma performance. These events are sudden and highly distinguishable from the frequency behaviour at times when ALEs do not occur, and are typically observed in the upper part of the kink/tearing/fishbone (KTF) frequency band (~ 1 kHz to 30 kHz) and the lower part of the toroidal Alfvén eigenmode (TAE) frequency band (~ 50 kHz to 200 kHz).[4]

Tokamaks feature a high-dimensional parameter space (ion temperature, magnetic flux density, ion number density, etc.) with large operational domains. For certain plasma parameters, stability transition has been shown to suddenly occur over small regions of parameter space, as is observed with the L-H mode transition[5] and edge localized mode (ELM) crashes[6]; accordingly, certain regions of parameter space have boundaries between different states of plasma stability. However, the transition that leads to ALEs is not fully understood; previous work has shown correlation between fast ion beta and neutral beam injection (NBI) beam energy[7], but a large area of parameter space still requires analysis.
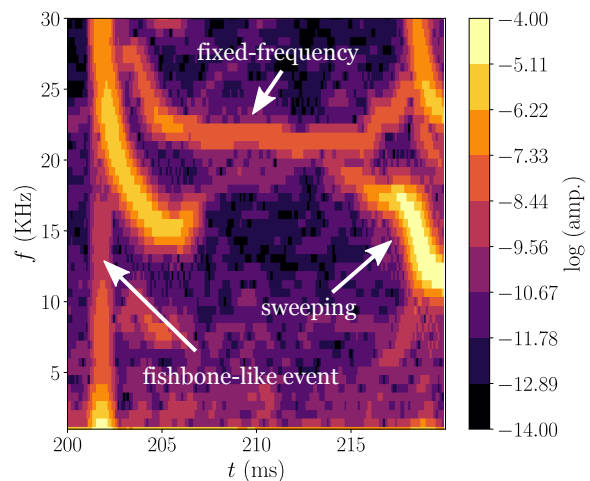
FIG. 1: Frequency spectrogram for magnetic fluctuations on NSTX in the 1 kHz to 30 kHz (kink/tearing/fishbone) band, 200 ms to 220 ms after the beginning of shot 139317. 3 types of mode character are observed: fixed-frequency eigenmodes, sweeping eigenmodes, and a fishbone-like event.

Previous work has highlighted the relationship between microturbulence, stochastic effects and fast ion loss[8,9], showing increased suppression of Alfvénic chirping as a function of microturbulence; theoretically microturbulence is treated as a stochastic term in the pitch-angle

scattering rate, while experimentally this can be heuristically inferred via the ion thermal conductivity χ_i. Additionally, work by Van Zeeland *et al.* has shown that Alfvénic chirping is increased in negative triangularity plasmas, where the microturbulence is expected to be lower.[10] However, there are other correlations which may yet be undiscovered.

Unfortunately, as one might expect in high-dimensional spaces, observing and predicting correlations becomes increasingly more difficult. With traditional computational analysis, there is no way to circumvent the amount of time required to perform human categorisation; a simple back-of-the-envelope calculation shows that even with an average characterisation time of 3 seconds per system state, the time taken to reach $\mathcal{O}(10^4)$ characterisations is over 8 hours. Previous work by Haskey *et al.* uses data mining techniques to to extract plasma fluctuations in time-series data, allowing one to identify different events using unsupervised classification.[11,12]

Machine learning (ML) also offers an increased capability to perform correlation studies in high-dimensional spaces. Here, we present correlation studies performed using NSTX data, and describe the newly developed ML framework ERICSON (Experimental Resonant Instability Correlation Studies on NSTX), examining plasma wave frequency-chirping observed via Mirnov coils. ERICSON allows us to compare NSTX data with a wide range of parameters that would be largely unfeasible with human classification, allowing us to understand which parameters affect EP transport. After initial training of the algorithm, we find the time taken for a single process to reach $\mathcal{O}(10^4)$ characterisations is under 10 minutes, allowing for one to examine a much higher number of characterisations, and therefore a much richer set of correlations. While ML is not asymptotically convergent to perfect accuracy, it allows for broad, statistical recognition of the plasma stability boundaries that do exist. One expects any erroneous characterisations to still allow for asymptotically correct stability boundaries as one tends to an infinite number of characterisations.

Magnetic perturbations in a very broadband range (0 to $\sim 500\,\text{MHz}$) are commonly measured on tokamaks such as NSTX by using Mirnov coils. Here, 5 characterisations of the frequency of magnetic perturbations are examined for KTF modes (of which 3 are shown in FIG. 1): quiescence (noise, or no frequency dependent behaviour), approximately fixed-frequency eigenmodes (herein referred to as simply fixed-frequency), sweeping eigenmodes (slow frequency variation due to time evolution of the plasma equilibrium), chirping (rapid frequency variation over a narrow frequency band due to wave-particle interaction), and fishbone-like (energetic particle mode with rapid frequency variation over a broad frequency band). For TAEs, we also examine 4 characterisations (see FIG. 2): quiescence, fixed-frequency, chirping, and ALEs (rapid frequency variation over a broad frequency band). Data from NSTX experiments

in 2010 is utilised, revealing a rich set of correlations between different mode character and weighted averages of plasma parameters obtained from TRANSP[13]. Our results are in agreement with previous work by Fredrickson et al.[7] using human characterisation in a reduced parameter space, allowing us to heuristically confirm the validity of predictions made by ERICSON. We show new, strong correlations between moments of the spectrograms and mode character, as well as evidence of a TAE stability boundary along $v_\varphi \approx \frac{1}{4}(v_{\text{inj.}} - 3v_\text{A})$.

II. FAST ION LOSS

Fast ions carry significantly more energy per mole than the thermal ions; on NSTX the NBI peak energy ($U_{\text{inj.}}$) lies at around $90\,\text{keV}$ and the thermal peak lies at $\sim 1\,\text{keV}$ to $2\,\text{keV}$ during typical tokamak operation[14]. The overall plasma performance can be heavily reduced in cases of large fast ion losses.[15] Fast ions give a significant contribution to the plasma pressure and are essential to transfer heat to the thermal population.

During frequency chirping events, non-linear structures known as 'holes' and 'clumps' can form on the ion distribution function, existing respectively as a relative decrease and increase of the local distribution function.[16–19] As the hole and clump move along the distribution function, they drag the resonant wave frequency with them. The canonical momentum space drift leads to a kinetic transport process; as a particle decreases in p_φ, it moves to a flux surface at higher minor radius.

NBI heating and RF heating can lead to a quasi-steady 'slowing-down' distribution of the form[20,21]:

$$F_{Si}(v) = \frac{3}{4\pi \ln(1 + v_\text{c}^3/v_{\text{inj.}}^3)} \frac{n_i}{v_\text{c}^3 + v^3} H(v_{\text{inj.}} - v) \quad (1)$$

where F_{Si} is the slowing-down distribution function for the i^th fast ion species, H is the Heaviside step function, n_i is the equilibrium number density of the i^th fast ion species, and v_c is the so-called critical speed:

$$v_\text{c} \equiv v_\text{te} \left(\frac{3\sqrt{\pi}}{4} \sum_i \frac{n_i m_\text{e}}{n_\text{e} m_i} Z_i^2 \right)^{1/3} \quad (2)$$

where v_te is the electron thermal speed, n_e is the equilibrium electron number density, m_i is the mass of the i^th fast ion species, m_e is the electron rest mass, and Z_i is the atomic number of the i^th fast ion species. Here, quasi-steady refers to the approximation that the distribution evolves smoothly as a function of $P_{\text{inj.}}$; fundamentally, this assumes that the rate of increase of injected power $1/\tau_\text{ramp}$ is much less than the reciprocal of the Spitzer slowing down time.
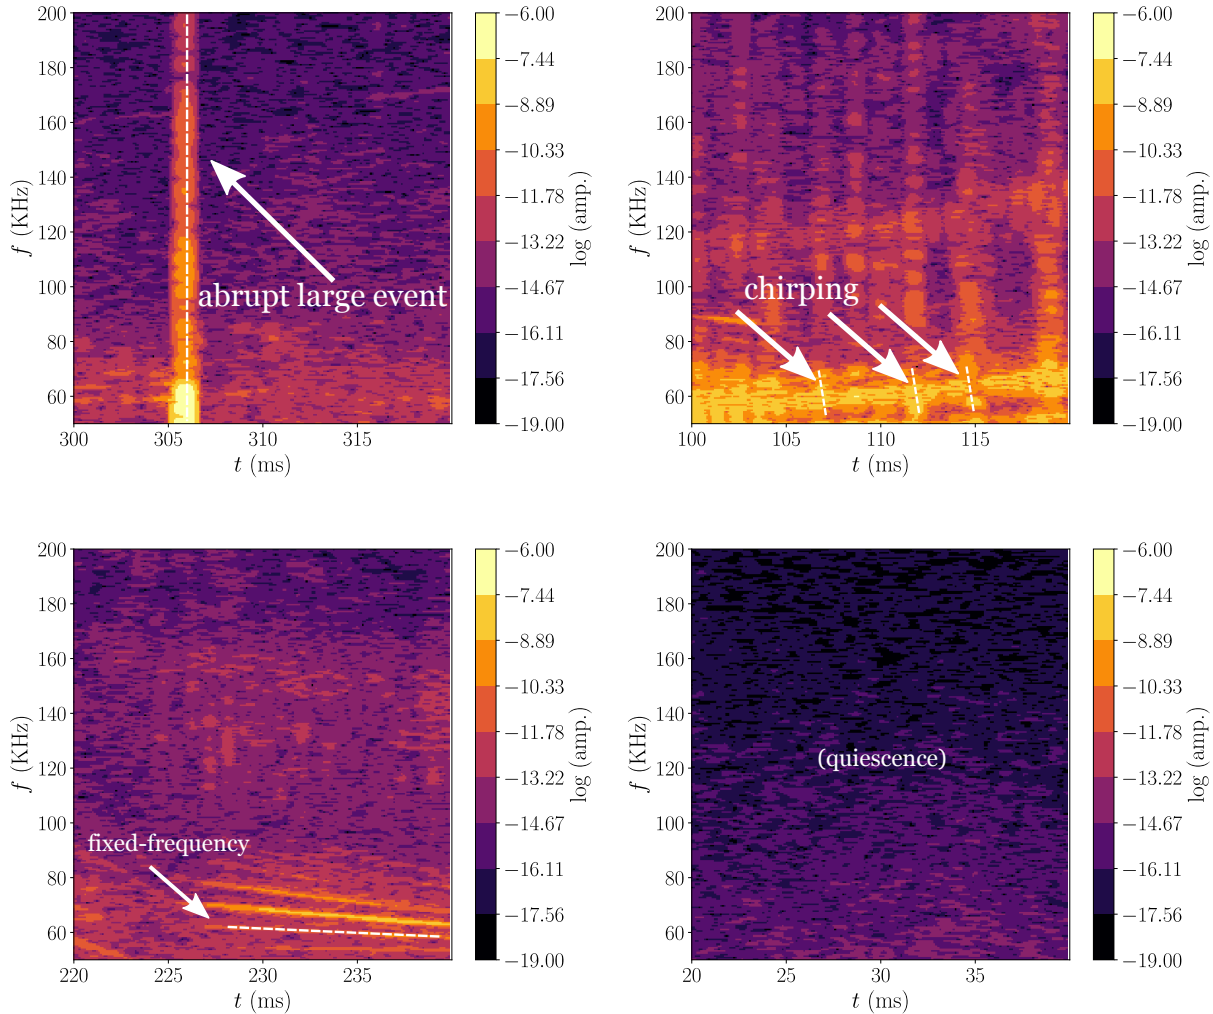
FIG. 2: Plots showing the 4 characterisations of toroidal Alfvénic eigenmode (TAE) magnetic fluctuations obtained from shot 139317 on NSTX. Top left: abrupt large event, top right: chirping, bottom-left: fixed-frequency, bottom-right: quiescence.

In the presence of background dissipation, kinetic instabilities that lead to hole and clump formation are triggered by gradients in the toroidal canonical angular momentum in a manner akin to inverse Landau damping; slowing distributions feature large gradients in the neighbourhood of $v_{\text{inj.}} \equiv \sqrt{2U_{\text{inj.}}/m}$, and therefore wave-particle interaction is enabled for waves with a resonance near $v_{\text{inj.}}$. As one attempts to increase the core temperature of the plasma, one creates sharper momentum gradients, leading to greater plasma instability. On NSTX, the NBI beam energy is super-Alfvénic ($v_{\text{inj.}} > v_{\text{A}}$), which allows for increased kinetic instability of Alfvén waves.

The formation of gap TAEs (gap toroidal Alfvén eigenmodes) allow for relatively long lifetime waves; these waves are not dispersive in radius, and therefore allow for long range frequency chirping in the $\sim 50\,\text{kHz}$ to $200\,\text{kHz}$

range[1]; the larger the frequency chirp, the further the momentum drift of the resonant particles.

Furthermore, experiments and simulations in the literature[2,15,22,23] have shown that rapid long range frequency chirping across a wide range of frequency chirps (mode 'avalanching') is correlated with high amounts of fast ion loss. Other work has shown that mode-mode destabilisation may play a role in Alfvénic frequency chirping - activity in the KTF frequency range ($\sim 1\,\text{kHz}$ to $30\,\text{kHz}$) may instigate Alfvénic activity, and vice-versa.

Crucially, the conditions which trigger small scale frequency chirping and so-called mode avalanching are not fully understood. Intuitive knowledge related to momentum gradients provides insight, but factors such as the q-profile, energy confinement time, and the fast ion density gradient may also play a role.

KTF character	Frequency traces
● (quiescent)	Noise
● (fixed-frequency)	Constant $\omega(t)$
● (sweeping)	Varying $\omega(t)$ ($\gg 1\,\text{ms}$)
● (chirping)	Rapidly varying $\omega(t)$ ($\sim 1\,\text{ms}$)
● (fishbone-like)	Rapidly varying $\omega(t)$ (broadband, $\sim 1\,\text{ms}$)

a)

TAE character	Frequency traces
● (quiescent)	Noise
● (fixed-frequency)	Constant $\omega(t)$
● (chirping)	Rapidly varying $\omega(t)$ ($\sim 1\,\text{ms}$)
● (ALEs)	Rapidly varying $\omega(t)$ (broadband, $\sim 1\,\text{ms}$)

b)

FIG. 3: a): Table detailing the characterisations of the mode frequency explored for kink/tearing/fishbones. The traces here refer to continuous streaks of slowly changing colour as one progresses in time on $\hat{B}(f,t)$ spectrograms. b): Table detailing the characterisations of the mode frequency explored for TAEs.

III. FRAMEWORK

ERICSON utilises four key parts: pre-processing NSTX data, mode characterisation, parameter space tracking from TRANSP data, correlation studies. Here, we discuss the pre-processing, mode characterisation, and the parameter space tracking; later, we show results from some correlation studies.

A. Pre-processing (NSTX data)

The induced voltage was measured across Mirnov coils. Via the Faraday-Lenz law:

$$V \equiv \oint_{\partial A} \mathbf{E} \cdot d\ell = -\partial_t \int_A \mathbf{B} \cdot d\mathbf{A} \qquad (3)$$

where V is the induced voltage, ∂A is the boundary of the area A of the coil, $\mathbf{E}$ is the electric field, $\mathbf{B}$ is the magnetic flux density, and $\hat{\mathbf{a}}$ is the unit vector normal to the surface, such that $d\mathbf{A} = \hat{\mathbf{a}}\,dA$.

We utilise the following asymmetric definitions for the discrete Fourier transform:

$$\mathcal{F}[B](\mathbf{r}, f[j]) = \sum_{l=0}^{N_t-1} B(\mathbf{r}, t[j]) \exp\left(i2\pi f[j]t[l]\right) \qquad (4)$$

$$\mathcal{F}^{-1}[\tilde{B}](\mathbf{r}, t[j]) = \frac{1}{L_t} \sum_{l=0}^{N_t-1} \tilde{B}(\mathbf{r}, f[l]) \exp\left(-i2\pi f[l]t[j]\right) \qquad (5)$$

where $\mathbf{r}$ is the position in real space, $f = \{l \in Z : 0 \leq l < N_t\}$ is a set of N_t frequencies, N_t is the number of time points in the dataset, and L_t is the temporal length of the dataset. Then, Fourier transforming the Faraday-Lenz law yields:

$$\tilde{V} \propto -i2\pi f\tilde{B} \qquad (6)$$

One can immediately see that $\tilde{B}$ is singular at $f = 0$. Because discrete analysis is used, this singularity becomes broadened and can affect nearby points. For this reason, we employ a cut off frequency ($f_{\text{min.}} \equiv 1\,\text{kHz}$) to avoid the singular value and nearby points saturating the data set. To minimise pick-up from the error-field correction switching power amplifiers (SPAs), we take the average time-domain signal from two Mirnov coils in close proximity of each other. However, one may instead choose to produce spectrograms from the cross-correlation $S_{\text{cr.}}$ of the two signals $\{V_1, V_2\}$ such that:

$$S_{\text{cr.}}(t) = \int_{-\infty}^{\infty} V_1(\tau)V_2(\tau + t)\,d\tau \qquad (7)$$

Under Fourier transform, this yields the product of the Fourier transform of the two signals. Here, we assume that the variation in signal amplitude and phase is small between the two coils, such that:

$$|\tilde{V}(\omega)| = \left|\frac{\tilde{V}_1 + \tilde{V}_2}{2}\right| \approx \sqrt{|\tilde{S}_{\text{cr.}}(\omega)|} \qquad (8)$$

As a result, for cases where the signal amplitude or phase varies appreciably between the two coils, it may prove advantageous to use the cross-correlation instead (for further details, please see Appendix A). To enable analysis of time-dependent frequencies, we use a short-time Fourier transform (STFT) to track the frequency evolution of modes in the system. The frequency resolution Δf and maximum frequency $f_{\text{max.}}$ obtained via an STFT are given respectively by:

$$\Delta f = \frac{1}{L_t}; \qquad f_{\text{max.}} = \frac{1}{2\Delta t} \qquad (9)$$

where Δt is the time resolution of the data set, and the data set is now more specifically defined as the time points within the STFT window. The forward-difference STFT is calculated as follows:

$$\hat{V}(f,t) = \mathcal{F}\left[\{W(t')V(t') : t < t' < t + L_t\}\right] \qquad (10)$$

where t is the time at which the STFT begins, $t'[j] = j\Delta t'$ is a dummy time used for the transform, $W(t')$ is

a window function, and $V(t)$ is now the two Mirnov coil average of the induced voltage. Many window functions are employed in the literature[24], each producing spurious sidebanding for a single frequency input. Here we employ a Hanning window due to its favourable decibel tapering for the erroneous signal produced, and its simple form:

$$W(t') = \sin^2\left(\frac{\pi t'}{N_t - 1}\right) \tag{11}$$

Using this method, one obtains the spectrogram signal:

$$\hat{B}(f,t) \propto \frac{i}{2\pi f}\hat{V}(f,t) \tag{12}$$

Finally, one can separate the spectrogram into different subdomains of f, allowing for characterisation of the signal in different frequency bands. A human would typically examine spectrograms obtained from this preprocessing, and perform characterisation (see FIG. 2 and FIG. 3 for characterisations).

For ML training and analysis, 125 shots from the 2010 NSTX archives were selected due to the clear observation of KTF and TAE activity. The Mirnov coils produce a signal at sampling rate $f_V = 5\,\mathrm{MHz}$ (200 ns resolution).

The shots were split into 20 ms slices, allowing for a maximum of ~ 20 chirping events; typically chirping in these frequency bands occurs on a $\sim 1\,\mathrm{ms}$ timescale.

The STFT time window L_t contained 2^{13} samples (1.6384 ms), allowing for a frequency resolution of $\sim 0.610\,\mathrm{kHz}$. By sliding the STFT using increments of 2^9 samples, we were able to obtain a time resolution τ on the spectrogram of 102.4 μs. It is instructive to note that by using a sliding increment smaller than the STFT window, to produce overlaid spectrograms that are plotted here (such as in FIG. 2) one must add a 'lag' of 819.2 ns as the pixel is actually the pixel at the center of the STFT window.

The 2D data for each slice was then converted into a contiguous 1D array in memory, made by taking contiguous blocks in frequency of the original spectrogram. That is to say:

$$\mathrm{ad}.[\hat{B}(f_2,t)] >> \mathrm{sizeof}(\hat{B}) = \mathrm{ad}.[\hat{B}(f_1, t+\tau)]$$
$$\forall f < f_2 : \mathrm{ad}.[\hat{B}(f,t)] >> \mathrm{sizeof}(\hat{B}) = \mathrm{ad}.[\hat{B}(f+\Delta f,t)]$$

where for a given spectrogram, f_1 is the minimum frequency, f_2 is the maximum frequency, $\mathrm{ad}.[\hat{B}(f,t)]$ is the address in memory of the spectrogram signal at (f,t), and $>>$ is the incremental bit-shift operator. A bit-shift by the datasize of the signal (sizeof($\hat{B}$)) allows us to relate two neighbouring array elements.

In total, a database of 4773 slices was generated. It is worth noting that there is a finite probability of an exact overlap of the plasma parameters between two slices in the database, although this is quite unlikely.

B. ML algorithm and training

In this paper, we employ supervised machine learning to characterise the spectral data. We have used random forest classifiers (RFCs), an ensemble variation of the well-established decision tree classifier (DTC). RFCs have been utilised and are well described in the literature[25,26] allowing for a simple, white-box approach of the problem, something we believe to be important for correlation studies.

DTCs act as flowchart-like algorithms which are optimised using a greedy algorithm. Therefore, while these classifiers are easy to operate, they sometimes only find local minima in the Gini impurity, leading to erroneous categorisation. Increasing the complexity of the tree (adding depth to the decision tree) can sometimes 'kick' the algorithm from a local minima into a global minima and improve accuracy, but too much depth leads to overfitting.

For this reason, we use RFCs which initialise an ensemble of decision trees with different random initial states. Then, the RFC returns the average probability that the data is of a given classification. This leads to a $\mathcal{O}(N_{\mathrm{DTC}})$ increase in accuracy for small N_{DTC}, where N_{DTC} is the number of trees.

Here, we examine a very high dimensional feature space (each spectrogram pixel is a single dimension) with an extremely high amount of collinearity; mode character is observed by relating the spectrogram amplitude of local clusters of pixels. RFCs subdivide the feature space on a Cartesian grid, with hyperparameters determining the available number of subdivisions (tree depth) and the size of the minimum undivided volume (leaf size).

Unfortunately, highly collinear, high dimensional spaces are hard for RFCs to efficiently divide. Therefore we expect RFCs to yield a lower accuracy than other ML schemes such as convolutional neural networks (CNNs). However for correlation studies, we examine broad, collective behaviour in parameter space; as a result we do not require extremely high accuracies.

The RFCs here are implemented using the Python library `scikit-learn`[27]; it is feasible that one could perform similar analysis in `tensorflow`, or write an RFC from scratch.

Human classification was performed using 10 shots, producing 337 slices for the KTF band (1 kHz to 30 kHz) and 281 slices for the TAE band (50 kHz to 200 kHz). We utilised training slices which featured low ambiguity in the frequency behaviour; some of these slices included multiple characterisations during the same time slice.

It is worth noting that here ERICSON does not check the D_α signal (light emission detected from the Balmer-alpha line) against the classification. As such, ERICSON could in principle classify ELMy time slices as featuring ALEs. However, the fraction of time points with ELMs, F, is very small. In FIG. 4, we present F using a histogram given by binning across different shots:
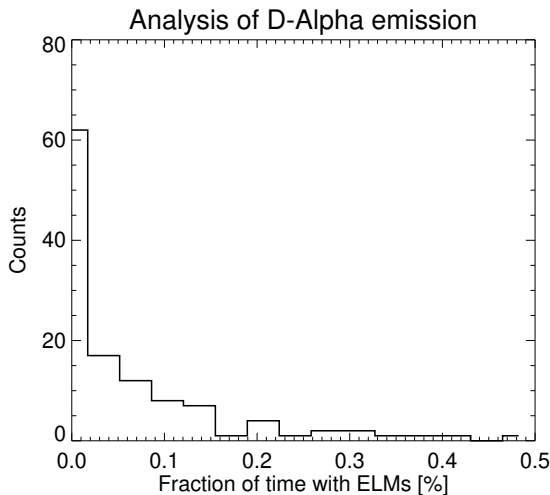
FIG. 4: Histogram for the number of shots with ELMs (edge localized modes) versus the fraction of time points in that shot with ELMs, using spikes in the D_α emission obtained from a fast ion D_α (FIDA) diagnostic. The shots featured here are the shots used in the correlation analysis performed in Section IV after training ERICSON. The fraction of time points containing ELMs is generally much less than 1% for the shots examined in this work, and as such we assume that any misclassification of ELMs as abrupt large events can be considered a negligible contribution to the correlations observed in Section IV.

$$F = \frac{n_{\mathrm{ELM}}}{f_D(t_1 - t_0)} \qquad (13)$$

where n_{ELM} is the number of time points with an ELM spike, f_D is the sampling frequency of the fast ion D_α diagnostic, and $t_1 - t_0$ is length of the time window of interest for each shot. In general, $F \ll 1\%$, and therefore we find that for the overall correlations given in Section IV, checking against the D_α signal was not deemed necessary for this work. However, for ELMy discharges, one would need to consider the D_α signal. We desire that characterisation falls under the following hierarchy for the KTF band:

fishbone-like → chirping → sweeping → fixed-freq. → quie.

with leftmost as the most important feature. That is to say, we desire a multi-class classifier such that any ambiguity leads to a more leftmost characterisation. For the TAE band, we desire the following hierarchy:

ALEs → chirping → fixed-freq. → quie.

with leftmost as the most important feature. A set of RFCs were then trained to perform binary classification

separately; each RFC yields a binary output for whether the system is exhibiting each of the types of character. We employ separate binary classification to allow us to tweak multi-class classification in a more direct fashion. The RFCs were trained using 75% of the learning set, and tested on the remaining 25%. The tree depth and number of trees were manually tweaked to optimise multi-class classification accuracy.
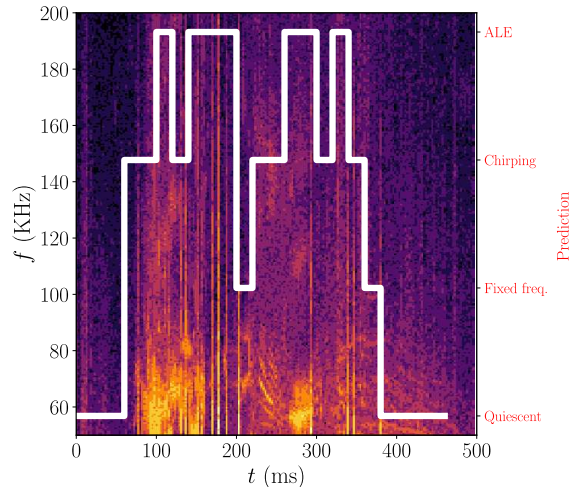

FIG. 5: Frequency spectrogram for magnetic fluctuations on NSTX in the 50 kHz to 200 kHz (toroidal Alfvén eigenmode) band for the first 500ms of shot 139317, overlaid with the mode character classified by ERICSON. 4 types of mode character are observed: quiescence, fixed-frequency eigenmodes, chirping modes, and abrupt large events (ALEs).

In FIG. 5, we show a frequency spectrogram for magnetic fluctuations found during NSTX shot 139317, overlaid with the classification made by ERICSON. As can be seen, ERICSON categorizes fairly cautiously, and sometimes characterises chirping behaviour as ALEs. To classify, a naïve approach might simply rely on using the most likely classification. In pseudocode one can write this as:

```
for (i = 0; i < K; i ++) {
    prob[i] = RFC(i)
}
return max_index(prob)
```

$$C_{\mathrm{KTF}}^{(\mathrm{train})} : \begin{bmatrix} 140 & 0 & 0 & 0 & 0 \\ 1 & 54 & 0 & 0 & 0 \\ 0 & 0 & 12 & 0 & 0 \\ 0 & 0 & 0 & 10 & 0 \\ 0 & 0 & 0 & 0 & 18 \end{bmatrix} \; ; \; C_{\mathrm{KTF}}^{(\mathrm{test})} : \begin{bmatrix} 57 & 2 & 1 & 0 & 4 \\ 1 & 16 & 3 & 1 & 0 \\ 2 & 0 & 2 & 0 & 2 \\ 0 & 0 & 0 & 1 & 1 \\ 2 & 0 & 0 & 0 & 8 \end{bmatrix} \qquad C_{\mathrm{TAE}}^{(\mathrm{train})} : \begin{bmatrix} 138 & 0 & 0 & 0 \\ 0 & 12 & 0 & 0 \\ 0 & 1 & 22 & 0 \\ 0 & 0 & 0 & 24 \end{bmatrix} \; ; \; C_{\mathrm{TAE}}^{(\mathrm{test})} : \begin{bmatrix} 59 & 0 & 0 & 1 \\ 1 & 2 & 1 & 2 \\ 0 & 2 & 7 & 1 \\ 0 & 0 & 3 & 6 \end{bmatrix}$$

Metric	Training set	Test set
Accuracy	99.2%	69.0%
R_K	0.996	0.824

Metric	Training set	Test set
Accuracy	98.9%	72.6%
R_K	0.995	0.871

FIG. 6: Confusion matrices and useful metrics from ERICSON trained to predict different types of classification for magnetic fluctuations on NSTX in the 1 kHz to 30 kHz (kink/tearing/fishbone) band and 50 kHz to 200 kHz (toroidal Alfvén eigenmode) band. From top-left to bottom-right, the leading diagonal of the confusion matrix for KTFs corresponds to correct classification for quiescence, fixed-frequency, frequency sweeping, chirping, and fishbone-like. From top-left to bottom-right, the leading diagonal of the confusion matrix for TAEs corresponds to correct classification for quiescence, fixed-frequency, chirping, and abrupt large events. The confusion matrices have high sparsity in the lower-diagonal part, owing to the preferential bias incorporated in ERICSON to increase prediction accuracy and enforce safer predictions.

where K is the number of binary classifiers used. This is typically the method employed by most multi-class classifiers. The classification accuracy from this method is typically $\min(\{p_i\})$ at best where p_i is the accuracy of the i^{th} RFC, making it a reasonable approach. However, due to the desired hierarchy, we instead impose a heirarchal method:

```
for (i = 0; i < K; i++) {
   if RFC(i) > RFC(l) + tol(i) for all l > i {
     return i
   }
}
```

Here, `tol(i)` denotes a tolerance factor. The confusion matrix $\{C_{ij}\}$ is defined here as the number of predictions of class j that were actually of class i, and the tolerance factor for each characterisation is therefore prescribed such that the confusion matrices for the training and test set meet two constraints: that the confusion matrices are close to upper-triangular, and that the confusion matrices are as close to diagonal as possible. The latter is the most important constraint, as a diagonal confusion matrix denotes perfect characterisation. By using the tolerance factor, we can force ERICSON to make better decisions by making it 'more cautious'; ERICSON would then preferentially classify TAE activity in a slice as chirping when previously classified as quiescent, if the tolerance is set to be higher. For moderate tolerance levels, this reduces overall confusion between different characterisations, but naturally creates some off-diagonal elements in the above diagonal part of the confusion matrices. This is deemed to be an acceptable error as we have encoded a preference for caution; we would rather have ERICSON predict a more 'dangerous' behaviour if it does misclassify.

In general, many of the characterised slices will have multiple features. Unfortunately this means that while the 'most likely' method produces the best accuracy out of the multi-class algorithms we employed, it has a finite ceiling on the accuracy, owed to the fact that it will not adequately distinguish between two features which appear at the same time during a slice. While the ML multi-class classifier could be improved, we find the accuracy to still be acceptable. In kind, a human performing classification could quite easily misinterpret a slice; what may be recorded as slow frequency sweeping by one human could be recorded as fixed-frequency by another. For this reason, we include in FIG. 6 the more conventional metric of accuracy, as well as the R_K coefficient. This coefficient has an upper limit of 1, and a lower limit which is greater than or equal to -1 depending on the dataset[28]:

$$R_K = \frac{\sum_{klm} C_{kk} C_{lm} - C_{kl} C_{mk}}{\sqrt{\sum_k \left(\sum_l C_{kl}\right)\left(\sum_{l',k'=k} C_{k'l'}\right)} \sqrt{\sum_k \left(\sum_l C_{lk}\right)\left(\sum_{l',k'=k} C_{l'k'}\right)}} \tag{14}$$

The R_K coefficient is a generalisation of the Matthews Correlation Coefficient for multi-class classification. It is intrinsically related to the Pearson correlation coefficient, as it is directly expressed as the ratio of the covariance of

two K-dimensional data sets corresponding to predicted classifications ($\underline{\mathbf{X}}$) and actual classifications ($\underline{\mathbf{Y}}$), and the product of the standard deviations of those two data sets:

$$R_K = \frac{\text{cov}(\underline{\mathbf{X}}, \underline{\mathbf{Y}})}{\sqrt{\text{cov}(\underline{\mathbf{X}}, \underline{\mathbf{X}}) \cdot \text{cov}(\underline{\mathbf{Y}}, \underline{\mathbf{Y}})}} \tag{15}$$

where $\text{cov}(a, b)$ is the covariance of a and b. Because an unbalanced data set is used for training, we incur a reasonable amount of accuracy bias; the real data set is also unbalanced (for example, ALEs are much less common than periods of quiescence). While the R_K is not as intuitive as accuracy, this is a better measure of the quality of the ML classifier, as it weights the metric depending on this accuracy bias, and gives a metric whose value is resilient under the use of an unbalanced real data set.

C. Parameter space tracking (TRANSP data)

For each discharge, plasma parameters are obtained from time-dependent simulations through the tokamak transport code TRANSP[13,29].

Simulations use measured profiles of electron and ion density and temperature, toroidal rotation, impurity content (assuming carbon is the dominant impurity).[30,31]

The plasma equilibrium as a function of time is computed through the EFIT code[32] using measured coils and plasma currents, data from magnetic sensors and kinetic profiles as main constraints. When available, the q-profile measured using the motional Stark effect diagnostic (MSE) at the mid plane[33] is also used as additional constraint.

Simulations include modeling of the fast ion distribution evolution in time through the Monte Carlo NUBEAM module of TRANSP[34,35]. NUBEAM includes classical phenomena such as collisional Coulomb scattering and slowing-down, and atomic physics (e.g. charge-exchange and neutralization reactions). Input parameters such as neutral beam injection energy, injected current, and neutral beam waveforms are taken from the experiment.

The main sources of uncertainty in TRANSP simulations are experimental uncertainties in the measured plasma profiles and inaccuracy in the equilibrium reconstruction. Typical uncertainties in the measured thermal plasma profiles are within 10% or better, which would result in small changes of derived quantities such as Alfvén velocity and β values reported in this work.

Similarly, uncertainties in the equilibrium reconstruction would cause a scatter of a few % in the reported values. For the equilibrium, the main cause of uncertainty is the lack of constraints on the q-profile in the core when MSE data are not available.

As one performs more classification, one populates the parameter space with more points. One expects that for phase transitions occurring over small, finite regions of parameter space, ERICSON should reveal regions with largely the same classification. Previous work has shown that frequency chirps exist for stochastic lifetimes[9], with other work showing that stochastic transport mechanisms such as ion microturbulence can effect the likelihood of chirping and therefore the characterisations observed[8].

For these reasons, we look at the overall behaviour of KTF and TAE magnetic activity in different regions of parameter space; one should not fully trust each individual classification, but rather the overall behaviour and the location of characterisation boundaries.

Each pixel shown in the figures in Section IV is binned such that the value of the pixel is the most frequent characterisation in that bin.

IV. CORRELATIONS

A. Mode-weighted averaging

Chirping is a non-linear phenomenon, requiring wave-particle nonlinearity. If one aims to correlate chirping with plasma parameters, some of these parameters may be spatially dependent; accordingly the mode structure plays an important role.

Here, we model the spatial distribution of physical quantities dependent on Alfvénic activity in the system via Bayesian inference, such that physical quantities are conditionally distributed *given the mode structure*. As such, by modelling the mode structure by using a normalized Gaussian (and therefore taking the prior to be Gaussian), quantities dependent on Alfvénic activity can be determined via a posterior distribution which is also Gaussian.

Accordingly, one can construct conditional expectations of the plasma quantities, where integrals are weighted with a normalised Gaussian, given here by $w(\rho)$:

$$\langle g \rangle \equiv \int\limits_0^1 (g \cdot w) \, \mathrm{d}\rho \tag{16}$$

where $\rho = \sqrt{\Psi/\Psi_0}$, Ψ_0 is the magnetic flux at the last closed flux surface, g is the quantity to be spatially averaged, and $\langle g \rangle$ is the mode-weighted average.

Experimentally, the mode structure can be observed via reflectometry data, however this data cannot probe hollow density profiles. Unfortunately, for a sizable number of chirping cases, beam deposition and other effects can lead to hollow density profiles, preventing inference of the mode structure via reflectometry.

In lieu of fully reliable values for every shot analysed, we make three assumptions. First, one expects Alfvén waves to have a fairly narrow mode structure, and as such we approximate the standard deviation to be ~ 0.1.

(a) Normalized injection velocity versus normalized beam ion beta for kink/tearing/fishbones. Fixed-frequency modes largely confined to super-Alfvénic injection velocity ($v_{\text{inj.}} \gtrsim 2v_{\text{A}}$) and low beam ion beta ($\beta_{\text{beam},i} \lesssim 0.25\beta$).

(b) Normalized injection velocity versus normalized beam ion beta for TAEs. Quiescent behaviour largely confined to low beam ion beta ($\beta_{\text{beam},i} \lesssim 0.2\beta$).

(c) Normalized injection velocity versus normalized plasma toroidal velocity for kink/tearing/fishbones. Quiescent behaviour largely confined to sub-Alfvénic toroidal velocity ($v_{\varphi} \lesssim 0.05v_{\text{A}}$). Fixed frequency behaviour dominates for $v_{\varphi} \gtrsim 0.05v_{\text{A}}$.

(d) Normalized injection velocity versus normalized plasma toroidal velocity for TAEs. Quiescent behaviour largely confined to relatively low plasma rotation $\left(v_{\varphi} \lesssim \frac{1}{4}(v_{\text{inj.}} - 3v_{\text{A}})\right)$.
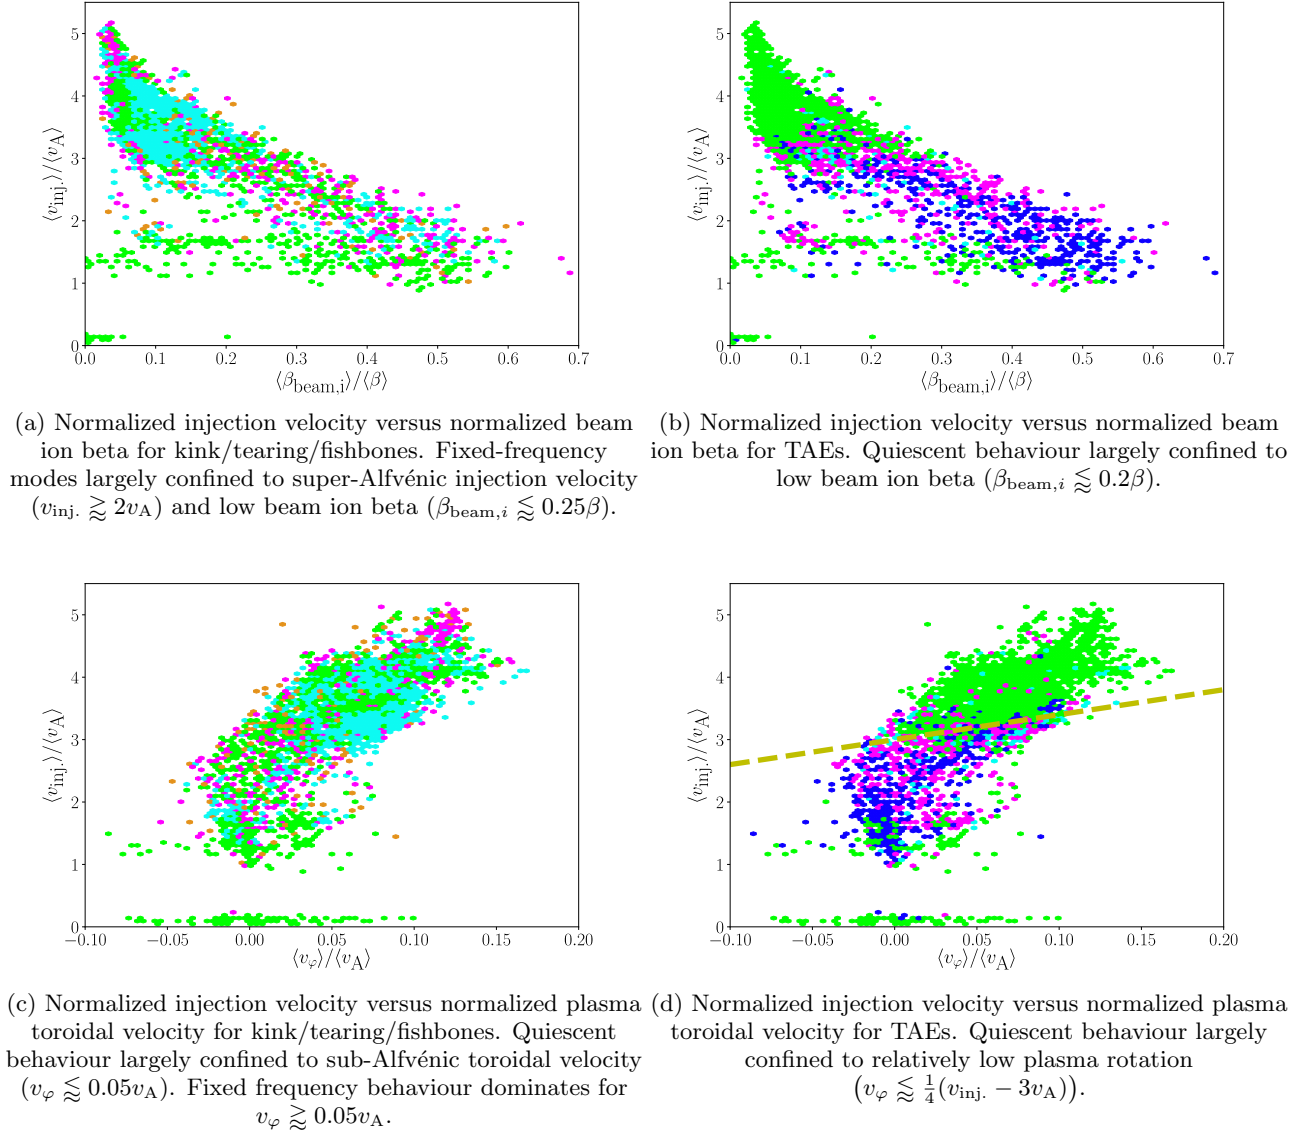
FIG. 7: Plots showing differing mode character as a function of operational parameters at NSTX. **Left plots:** kink/tearing/fishbones (1 kHz to 30 kHz modes); quiescent (green), fixed-frequency (cyan), sweeping (orange), chirping (blue), fishbone-like (magenta). **Right plots:** TAEs (50 kHz to 200 kHz modes); quiescent (green), fixed-frequency (cyan), chirping (blue), ALEs (magenta).

Second, we expect Alfvénic mode structure to peak at around q_{min},[1,36,37] which we approximate to be at $\rho = 0.5$ (see Levinton and Yuh[33]). Third, we extend this assumption also to the KTF modes in this analysis, however one would in reality expect the mode structure to be broader.[38] Therefore, we use the following trial function for analysis:

$$w(\rho) \propto \exp\left[-\left(\frac{(\rho - \sqrt{0.5})^2}{2(0.1^2)}\right)\right] \qquad (17)$$

such that the mode structure is approximated by a Gaussian with peak at $\rho = 0.5$ and standard deviation 0.1.

B. Injection velocity

FIG. 7 contains 4 plots showing differing mode character as a function of operational parameters at NSTX.
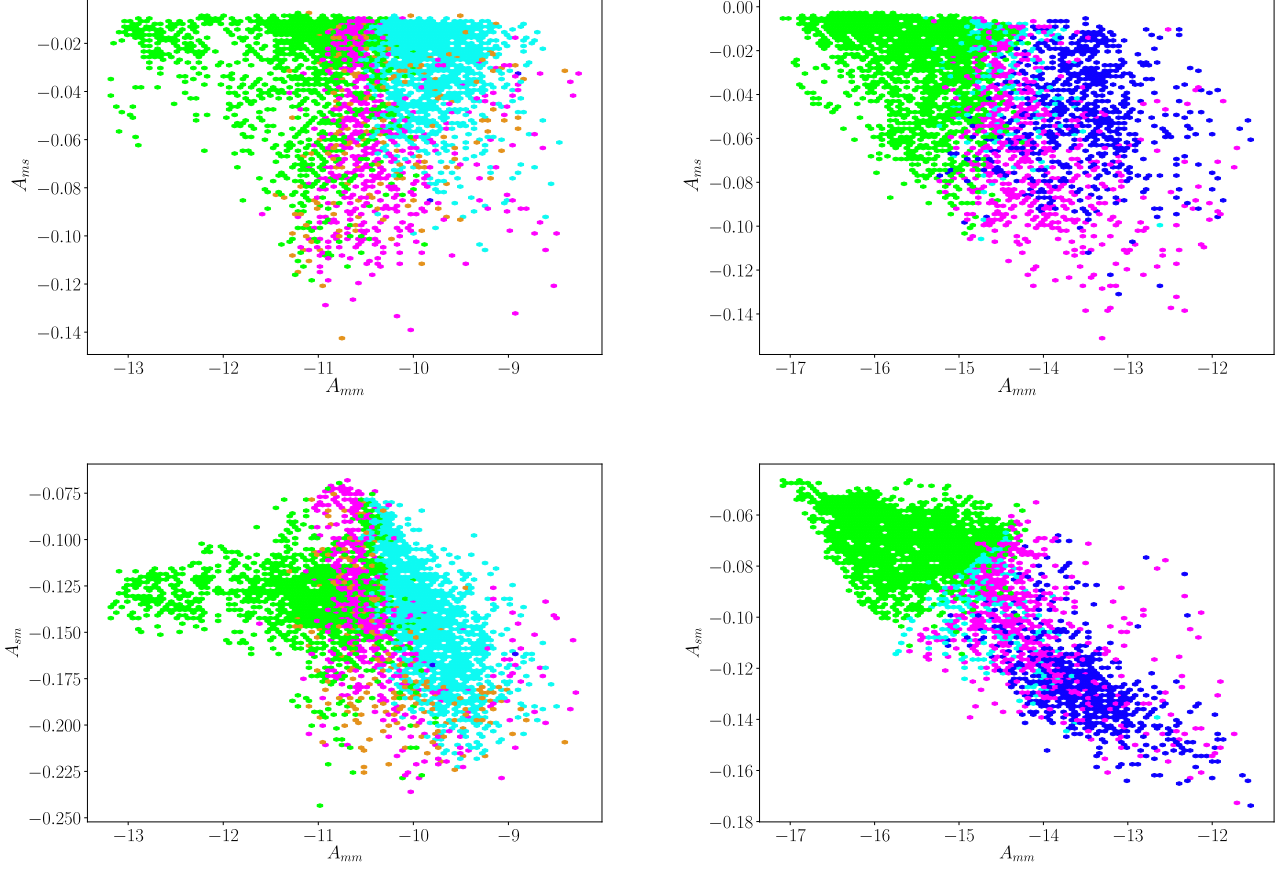
FIG. 8: Plots showing differing mode character as a function of spectrogram moments at NSTX. **Left plots:** kink/tearing/fishbones (1 kHz to 30 kHz modes); quiescent (green), fixed-frequency (cyan), sweeping (orange), chirping (blue), fishbone-like (magenta). Spectrogram average (A_{mm}) is a good indicator of mode character; average frequency spread (A_{sm}) and temporal intermittency (A_{ms}) are poor indicators. **Right plots:** TAEs (50 kHz to 200 kHz modes); quiescent (green), fixed-frequency (cyan), chirping (blue), ALEs (magenta). Spectrogram average (A_{mm}) and average frequency spread (A_{sm}) are good indicators of mode character; temporal intermittency (A_{ms}) is a poor indicator.

On the vertical axis for each plot is the normalized injection velocity $\langle v_{\mathrm{inj.}}\rangle/\langle v_{\mathrm{A}}\rangle$, where the Alfvén speed is given by:

$$\langle v_{\mathrm{A}}\rangle(t) \approx \frac{\langle|B|\rangle}{\sqrt{2\mu_0\langle n_e\rangle m_p}} \qquad (18)$$

where B is the magnetic flux density, n_e is the electron density, and m_p is the proton rest mass. This approximation assumes roughly 2 atomic mass units per electron, which is suitable for the NSTX plasma in these shots (typically featuring deuterium and carbon ions). The normalized beam ion beta is defined as $\langle\beta_{\mathrm{beam},i}\rangle/\langle\beta\rangle$, where $\langle\beta_{\mathrm{beam},i}\rangle$ is the beam ion beta, and $\langle\beta\rangle$ is the total beta.

The left hand plots in FIG. 7 show KTF mode character (1 kHz to 30 kHz modes); quiescent (green), fixed-

frequency (cyan), sweeping (orange), chirping (blue), fishbone-like (magenta). The right hand plots in FIG. 7 show TAE mode character (50 kHz to 200 kHz modes); quiescent (green), fixed-frequency (cyan), chirping (blue), ALEs (magenta).

For modes in the KTF band, fixed-frequency modes are largely confined to regions of parameter space where the injection velocity is super-Alfvénic, and the beam ion beta is relatively low. We find fixed-frequency mode behaviour for $v_{\mathrm{inj.}} \gtrapprox 2v_{\mathrm{A}}$, $\beta_{\mathrm{beam},i} \lessapprox 0.25\beta$, $v_\varphi \gtrapprox 0.05v_{\mathrm{A}}$. Quiescence is largely observed for sub-Alfvénic toroidal velocity, approximately given by $v_\varphi \lessapprox 0.05v_{\mathrm{A}}$.

For modes in the TAE band, quiescent behaviour is largely confined to regions of parameter space where the beam ion beta is relatively low. We find quiescence for $\beta_{\mathrm{beam},i} \lessapprox 0.2\beta$, $v_\varphi \lessapprox \frac{1}{4}(v_{\mathrm{inj.}} - 3v_{\mathrm{A}})$.
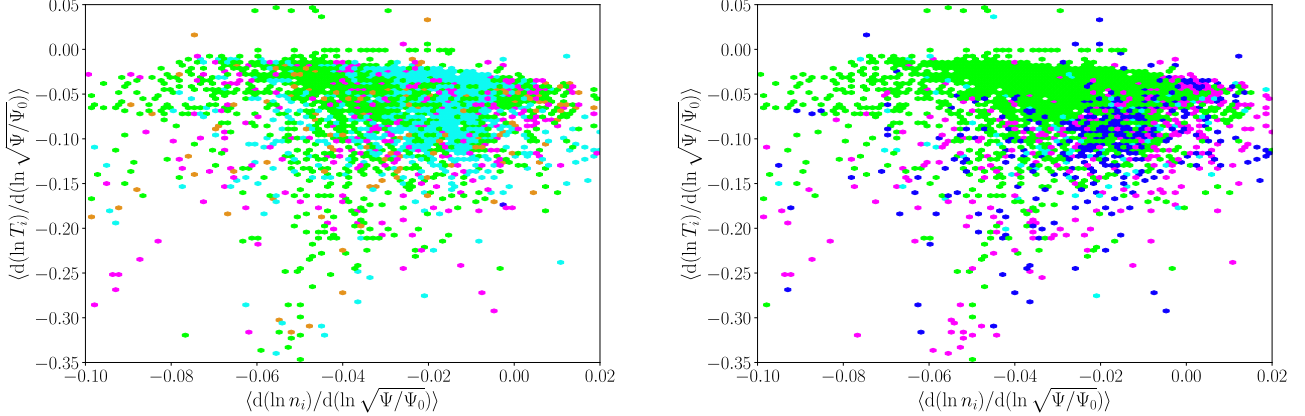
FIG. 9: Plots showing differing mode character as a function of normalized ion temperature gradient $(\mathrm{d}(\ln T_i)/\mathrm{d}(\ln \sqrt{\Psi/\Psi_0}))$ and normalized ion density gradient $(\mathrm{d}(\ln n_i)/\mathrm{d}(\ln \sqrt{\Psi/\Psi_0}))$ at NSTX. **Left plot:** kink/tearing/fishbones (1 kHz to 30 kHz modes); quiescent (green), fixed-frequency (cyan), sweeping (orange), chirping (blue), fishbone-like (magenta). Largely null result; no strong correlations observed. **Right plot:** TAEs (50 kHz to 200 kHz modes); quiescent (green), fixed-frequency (cyan), chirping (blue), ALEs (magenta). Chirping and ALEs occur at low $\eta = (\nabla T_i/\nabla n_i)$.

Plasma rotation can play a strong role in ideal MHD stabilisation; sub-Alfvénic plasma rotation serves to stabilise the plasma, leading to increased quiescence in the KTF frequency band. However, the TAEs are subject to kinetic instabilities for super Alfvénic ($v_{\mathrm{inj.}} > v_{\mathrm{A}}$) near the Alfvén speed. Our results allow us to further posit for Alfvén waves that if the injection velocity is less than the Alfvén speed, one expects that a reversed toroidal plasma velocity would lead to decreased TAE activity.

The beam ion beta also plays a strong role in TAE destabilisation. Increased NBI power increases the kinetic drive for resonant modes; this leads to an increased likelihood for nonlinear wave-particle interaction and chirping[39]. As the gradient of the fast-ion distribution determines the kinetic stability of nearby TAEs in momentum space, one expects theoretically that increased ion beam beta leads to increased TAE activity. The results from ERICSON show increased chirping and ALE activity at high beam ion beta in agreement with theory and observation.

C. Spectrogram moments

FIG. 8 contains 4 plots showing differing mode character as a function of spectrogram moments at NSTX. The left hand plots show KTF mode character (1 kHz to 30 kHz modes); quiescent (green), fixed-frequency (cyan), sweeping (orange), chirping (blue), fishbone-like (magenta). The right hand plots show TAE mode character (50 kHz to 200 kHz modes); quiescent (green), fixed-frequency (cyan), chirping (blue), ALEs (magenta).

We define three simple metrics based on moments of

the spectrogram signal given by (12), taken both in frequency and time to yield a scalar. First, we define the spectrogram average A_{mm} as:

$$A_{mm} \propto \sum_l \sum_j \hat{B}(f[j], t[l]) \qquad (19)$$

A_{mm} is a measure of the amplitude of magnetic fluctuations in a given frequency band. The average frequency spread A_{sm} is defined as:

$$A_{sm} \propto \sum_l \sqrt{\sum_j \left(\hat{B}(f[j], t[l]) - \mu_f(t[l]) \right)^2} \qquad (20)$$

where $\mu_f(t[l]) \propto \sum_j \hat{B}(f[j], t[l])$ is the frequency average of $\hat{B}$. A_{sm} is a measure of how broadband the signal is. The intermittency A_{ms} is defined as:

$$A_{ms} \propto \sqrt{\sum_l \left(\left[\sum_j \hat{B}(f[j], t[l]) \right] - A_{mm} \right)^2} \qquad (21)$$

A_{ms} is a measure of how intermittent the signal is.

For modes in the KTF band, the spectrogram average heavily dominates over the other moments; the mode character is largely invariant as a function of the average frequency spread and intermittency. For modes in the TAE band, the mode character is largely invariant as a function of the intermittency.

D. Ion η

FIG. 9 contains 2 plots showing differing mode character as a function of ion temperature and pressure gradients at NSTX. The left hand plot shows KTF mode character (1 kHz to 30 kHz modes); quiescent (green), fixed-frequency (cyan), sweeping (orange), chirping (blue), fishbone-like (magenta). The right hand plot shows TAE mode character (50 kHz to 200 kHz modes); quiescent (green), fixed-frequency (cyan), chirping (blue), ALEs (magenta).

For modes in the KTF band, we observe a largely null result; no discernible correlation can be identified. However, for modes in the TAE band, chirping and ALEs occur at low $\eta = (\nabla T_i / \nabla n_i)$. This is consistent with observations in DIII-D[8], reduced nonlinear kinetic simulations[9] and nonlinear gyrokinetic simulations performed for NSTX experiments[40]. Higher η implies more drive for ion temperature gradient (ITG) modes and therefore more turbulent stochastisation of phase space. This leads to suppression of nonlinear structures, such as holes and clumps, that sustain chirping. We note, however, that each turbulent mode has its own threshold in η, hence no global threshold in η can be identified as a well-defined transition for the chirping/fixed-frequency boundary.

V. CONCLUSIONS

We have developed an ML framework to expedite the physics analysis of Alfvén waves at NSTX. We employed Random Forest Classifiers to study correlations between plasma parameters and the frequency response of Alfvén waves, which indicates the nature of fast ion losses.

One possible extension of this work could be towards higher frequency Alfvén waves, such as co- and counter-propagating global Alfvén waves (GAEs). Their frequencies have recently been shown to be highly affected by beam parameters, such as the ratio of the injection speed to the Alfvén speed and the central pitch[41], with no appreciable eigenstructure modifications. The ML framework described here can be a useful attempt to unveil possible hidden correlations in the observations.

Further work building on this analysis could aim to examine correlations in the full parameter space. Transitions such as the L-H mode transition have been partially explained in reduced parameter spaces (such as the 2D parameter space with magnetic shear $\hat{s}$ and normalised pressure gradient α)[42]. This transition is shown by projecting the full parameter space onto a plane, similar to the plots we show; in reality, this transition occurs over more exotic surfaces in parameter space which are topologically challenging to analyse. Furthermore, it is entirely feasible that a coordinate transform in the parameter space would change the topology of stability boundaries - one cannot predict *a priori* what is the most sensible representation of the parameter space for a given stability condition.

The work in this paper can be expanded to examine automatic dataset development (such as archival searches for shots with given mode behaviour), real-time feedback control, and real-time modeling. Recent work using ML has investigated real-time capable modeling of NSTX-U using neural networks[43]; by also employing classification data, it may be possible to further enhance the predictive capabilities of such a framework.

For both the TAE and KTF bands, very strong correlations are found between mode character and moments of the spectrograms of magnetic fluctuations in the plasma found in Section IV C. Accordingly, one could use these moments as features instead; one can expect that reduction to this three dimensional space yields higher accuracy at quicker computational speeds.

VI. ACKNOWLEDGMENTS

Appendix A: Signal cross-correlation

If one represents the signal $V(\mathbf{r}, t)$ using a Fourier transform:

$$V(\mathbf{r}, t) = \int\limits_0^\infty \tilde{V}(\mathbf{r}, \omega) e^{i\omega t}$$

the complex Fourier spectrum $\tilde{V}(\mathbf{r}, \omega)$ can be represented by using an amplitude and phase:

$$\tilde{V}(\mathbf{r}, \omega) = |\tilde{V}|(\mathbf{r}, \omega) e^{i\phi(\mathbf{r}, \omega)}$$

Then, let us define the shorthand $A(\mathbf{r} = \mathbf{r}_j) = A_j$. Then, for two signals $\{V_1, V_2\}$, we define:

$$\delta \tilde{V}(\mathbf{r}_1, \mathbf{r}_2, \omega) = \tilde{V}_2 - \tilde{V}_1$$

such that $\delta \tilde{V}$ is the difference in spectra. We then define the following power spectra:

$$\tilde{S}_{\mathrm{cr.}}(\mathbf{r}_1, \mathbf{r}_2, \omega) = \tilde{V}_1 \tilde{V}_2; \quad \tilde{S}_{\mathrm{avg.}}(\mathbf{r}_1, \mathbf{r}_2, \omega) = \left(\frac{\tilde{V}_1 + \tilde{V}_2}{2} \right)^2$$

where $\tilde{S}_{\mathrm{cr.}}$ is the Fourier transform of the cross-correlation, and $\tilde{S}_{\mathrm{avg.}}$ is the square of the Fourier transform of the average. It is trivial to show in the limit that $|\delta \tilde{V}| \ll |\tilde{V}_1|$, to first order in $\delta \tilde{V}$ the power spectra are identical. Consequently, by separating into amplitude and phase, with some abuse of notation one can define a correlation length spectrum $L_{\mathrm{cr.}}(\omega)$:

$$\forall |\mathbf{r}_2 - \mathbf{r}_1| \leq L_{\mathrm{cr.}}(\omega) : \begin{cases} |\delta|\tilde{V}|(\mathbf{r}_1, \mathbf{r}_2, \omega)| & \ll |\tilde{V}_1|, |\tilde{V}_2| \\ |\delta \phi(\mathbf{r}_1, \mathbf{r}_2, \omega)| & \ll \pi \end{cases}$$

This typically requires that the change in amplitude is small, but also that the two signals are suitably correlated in phase. As an example, if $V_2 = -V_1$, such that $|\delta \phi(\mathbf{r}_1, \mathbf{r}_2, \omega)| = \pi$, then $|\delta|\tilde{V}|| \ll |\tilde{V}_1|, |\tilde{V}_2|$ is satisfied, but:

$$\tilde{S}_{\mathrm{avg.}} = 0$$

giving a significant difference in power spectra, as the phase difference is large.

[1] W. W. Heidbrink. Basic physics of Alfvén instabilities driven by energetic particles in toroidally confined plasmas. *Phys. Plasmas*, 15(5):055501, may 2008.

[2] E. D. Fredrickson, N. N. Gorelenkov, R. E. Bell, J. E. Menard, A. L. Roquemore, S. Kubota, N. A. Crocker, and W. Peebles. Fast ion loss in a 'sea-of-TAE'. *Nucl. Fusion*, 46(10):S926–S932, oct 2006.

[3] A. Bierwage, K. Shinohara, Y. Todo, N. Aiba, M. Ishikawa, G. Matsunaga, M. Takechi, and M. Yagi. Simulations tackle abrupt massive migrations of energetic beam ions in a tokamak plasma. *Nat. Commun.*, 9(1):3282, dec 2018.

[4] M. Podestà, W. W. Heidbrink, D. Liu, E. Ruskov, R. E. Bell, D. S. Darrow, E. D. Fredrickson, N. N. Gorelenkov, G. J. Kramer, B. P. LeBlanc, S. S. Medley, A. L. Roquemore, N. A. Crocker, S. Kubota, and H. Yuh. Experimental studies on fast-ion transport by Alfvén wave avalanches on the National Spherical Torus Experiment. *Phys. Plasmas*, 16(5):056104, may 2009.

[5] S.-I. Itoh and K. Itoh. Model of L to H-Mode Transition in Tokamak. *Phys. Rev. Lett.*, 60(22):2276–2279, may 1988.

[6] J. W. Connor. Edge-localized modes - physics and theory. *Plasma Phys. Control. Fusion*, 40(5):531–542, may 1998.

[7] E. D. Fredrickson, N. N. Gorelenkov, M. Podesta, A. Bortolon, S. P. Gerhardt, R. E. Bell, A. Diallo, and B. LeBlanc. Parametric dependence of fast-ion transport events on the National Spherical Torus Experiment. *Nucl. Fusion*, 54(9):093007, sep 2014.

[8] V. N. Duarte, H. L. Berk, N. N. Gorelenkov, W. W. Heidbrink, G. J. Kramer, R. Nazikian, D. C. Pace, M. Podestà, B. J. Tobias, and M. A. Van Zeeland. Prediction of nonlinear evolution character of energetic-particle-driven instabilities. *Nucl. Fusion*, 57(5):054001, may 2017.

[9] B. J. Q. Woods, V. N. Duarte, A. J. De-Gol, N. N. Gorelenkov, and R. G. L. Vann. Stochastic effects on phase-space holes and clumps in kinetic systems near marginal stability. *Nucl. Fusion*, 58(8):082015, 2018.

[10] M. A. Van Zeeland, C. S. Collins, W. W. Heidbrink, M. E. Austin, X. D. Du, V. N. Duarte, A. Hyatt, G. Kramer, N. Gorelenkov, B. Grierson, D. Lin, A. Marinoni, G. McKee, C. Muscatello, C. Petty, C. Sung, K. E. Thome, M. Walker, and Y. B. Zhu. Alfvén eigenmodes and fast ion transport in negative triangularity DIII-D plasmas. *Nucl. Fusion*, 59(8):086028, aug 2019.

[11] S. R. Haskey, B. D. Blackwell, and D. G. Pretty. Clustering of periodic multichannel timeseries data with application to plasma fluctuations. *Comput. Phys. Commun.*, 185(6):1669–1680, 2014.

[12] S. R. Haskey, B. D. Blackwell, C. Nührenberg, A. Könies, J. Bertram, C. Michael, M. J. Hole, and J. Howard. Experiment-theory comparison for low frequency BAE modes in the strongly shaped H-1NF stellarator. *Plasma Phys. Control. Fusion*, 57(9), aug 2015.

[13] J. Breslau, M. Gorelenkova, F. Poli, J. Sachdev, and X. Yuan. TRANSP, jun 2018.

[14] E. J. Synakowski, M. G. Bell, R. E. Bell, T. Bigelow, M. Bitter, W. Blanchard, J. Boedo, C. Bourdelle, C. Bush, D. S. Darrow, P. C. Efthimion, E. D. Fredrickson, D. A. Gates, M. Gilmore, L. R. Grisham, J. C. Hosea, D. W. Johnson, R. Kaita, S. M. Kaye, S. Kubota, H. W. Kugel, B. P. LeBlanc, K. Lee, R. Maingi, J. Manickam, R. Maqueda, E. Mazzucato, S. S. Medley, J. Menard, D. Mueller, B. A. Nelson, C. Neumeyer, M. Ono, F. Paoletti, H. K. Park, S. F. Paul, Y.-K. M. Peng, C. K. Phillips, S. Ramakrishnan, R. Raman, A. L. Roquemore, A. Rosenberg, P. M. Ryan, S. A. Sabbagh, C. H. Skinner, V. Soukhanovskii, T. Stevenson, D. Stutman, D. W. Swain, G. Taylor, A. Von Halle, J. Wilgen, M. Williams, J. R. Wilson, S. J. Zweben, R. Akers, R. E. Barry, P. Beiersdorfer, J. M. Bialek, B. Blagojevic, P. T. Bonoli, R. Budny, M. D. Carter, C. S. Chang, J. Chrzanowski, W. Davis, B. Deng, E. J. Doyle, L. Dudek, J. Egedal, R. Ellis, J. R. Ferron, M. Finkenthal, J. Foley, E. Fredd, A. Glasser, T. Gibney, R. J. Goldston, R. Harvey, R. E. Hatcher, R. J. Hawryluk, W. Heidbrink, K. W. Hill, W. Houlberg, T. R. Jarboe, S. C. Jardin, H. Ji, M. Kalish, J. Lawrance, L. L. Lao, K. C. Lee, F. M. Levinton, N. C. Luhmann, R. Majeski, R. Marsala, D. Mastravito, T. K. Mau, B. McCormack, M. M. Menon, O. Mitarai, M. Nagata, N. Nishino, M. Okabayashi, G. Oliaro, D. Pacella, R. Parsells, T. Peebles, B. Peneflor, D. Piglowski, R. Pinsker, G. D. Porter, A. K. Ram, M. Redi, M. Rensink, G. Rewoldt, J. Robinson, P. Roney, M. Schaffer, K. Shaing, S. Shiraiwa, P. Sichta, D. Stotler, B. C. Stratton, Y. Takase, X. Tang, R. Vero, W. R. Wampler, G. A. Wurden, X. Q. Xu, J. G. Yang, L. Zeng, and W. Zhu. The national spherical torus experiment (NSTX) research programme and progress towards high beta, long pulse operating scenarios. *Nucl. Fusion*, 43(12):1653, dec 2003.

[15] D. S. Darrow, N. Crocker, E. D. Fredrickson, N. N. Gorelenkov, M. Gorelenkova, S. Kubota, S. S. Medley, M. Podestà, L. Shi, and R. B. White. Stochastic orbit loss of neutral beam ions from NSTX due to toroidal Alfvén eigenmode avalanches. *Nucl. Fusion*, 53(1):013009, jan 2013.

[16] H. L. Berk, B. N. Breizman, and M. Pekker. Numerical simulation of bump-on-tail instability with source and sink. *Phys. Plasmas*, 2(8):3007–3016, aug 1995.

[17] H. L. Berk, B. N. Breizman, and M. Pekker. Nonlinear Dynamics of a Driven Mode near Marginal Stability. *Phys. Rev. Lett.*, 76(8):1256–1259, feb 1996.

[18] H. L. Berk, B. N. Breizman, and N. V. Petviashvili. Spontaneous hole-clump pair creation in weakly unstable plasmas. *Phys. Lett. A*, 234(3):213–218, sep 1997.

[19] Y. Todo. Introduction to the interaction between energetic particles and Alfven eigenmodes in toroidal plasmas. *Rev. Mod. Plasma Phys.*, 3(1):1, dec 2019.

[20] J. G. Cordey. Effects of particle trapping on the slowing-down of fast ions in a toroidal plasma. *Nucl. Fusion*, 16(3):499, 1976.

[21] J. D. Gaffey. Energetic ion distribution resulting from neutral beam injection in tokamaks. *J. Plasma Phys.*, 16(2):149–169, oct 1976.

[22] G. J. Kramer, M. Podestà, C. Holcomb, L. Cui, N. N. Gorelenkov, B. Grierson, W. W. Heidbrink, R. Nazikian, W. Solomon, M. A. Van Zeeland, and Y. Zhu. Improving fast-ion confinement in high-performance discharges by suppressing Alfvén eigenmodes. *Nucl. Fusion*, 57(5):056024, may 2017.

[23] C. S. Collins, W. W. Heidbrink, M. Podestà, R. B. White, G. J. Kramer, D. C. Pace, C. C. Petty, L. Stagner, M. A. Van Zeeland, Y. B. Zhu, and the DIII-D Team. Phase-space dependent critical gradient behavior of fast-ion transport due to Alfvén eigenmodes. *Nucl. Fusion*, 57(8):086005, aug 2017.

[24] F.J. Harris. On the use of windows for harmonic analysis with the discrete Fourier transform. *Proc. IEEE*, 66(1):51–83, 1978.

[25] L. Breiman. Random Forests. *Mach. Learn.*, 45(1):5–32, 2001.

[26] C. Rea and R. S. Granetz. Exploratory Machine Learning Studies for Disruption Prediction Using Large Databases on DIII-D. *Fusion Sci. Technol.*, 74(1-2):89–100, aug 2018.

[27] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in {P}ython. *J. Mach. Learn. Res.*, 12:2825–2830, 2011.

[28] J. Gorodkin. Comparing two K-category assignments by a K-category correlation coefficient. *Comput. Biol. Chem.*, 28(5-6):367–374, dec 2004.

[29] R.J. Hawryluk. AN EMPIRICAL APPROACH TO TOKAMAK TRANSPORT. In *Phys. Plasmas Close to Thermonucl. Cond.*, pages 19–46. Elsevier, 1981.

[30] B. P. LeBlanc, R. E. Bell, D. W. Johnson, D. E. Hoffman, D. C. Long, and R. W. Palladino. Operation of the NSTX thomson scattering system. In *Rev. Sci. Instrum.*, volume 74, pages 1659–1662, mar 2003.

[31] R. E. Bell, R. Andre, S. M. Kaye, R. A. Kolesnikov, B. P. Leblanc, G. Rewoldt, W. X. Wang, and S. A. Sabbagh. Comparison of poloidal velocity measurements to neoclassical theory on the National Spherical Torus Experiment. *Phys. Plasmas*, 17(8), aug 2010.

[32] L. L. Lao, H. St John, R. D. Stambaugh, A. G. Kellman, and W. Pfeiffer. Reconstruction of current profile parameters and plasma shapes in tokamaks. *Nucl. Fusion*, 25(11):1611–1622, 1985.

[33] F. M. Levinton and H. Yuh. The motional Stark effect diagnostic on NSTX. In *Rev. Sci. Instrum.*, volume 79, 2008.

[34] R. J. Goldston, D. C. McCune, H. H. Towner, S. L. Davis, R. J. Hawryluk, and G. L. Schmidt. New techniques for calculating heat and particle source rates due to neutral beam injection in axisymmetric tokamaks. *J. Comput. Phys.*, 43(1):61–78, 1981.

[35] Alexei Pankin, Douglas McCune, Robert Andre, Glenn Bateman, and Arnold Kritz. The tokamak Monte Carlo fast ion module NUBEAM in the national transport code collaboration library. *Comput. Phys. Commun.*, 159(3):157–184, jun 2004.

[36] H. L. Berk, D. N. Borba, B. N. Breizman, S. D. Pinches, and S. E. Sharapov. Theoretical Interpretation of Alfvén Cascades in Tokamaks with Nonmonotonic q Profiles. *Phys. Rev. Lett.*, 87(18):185002, oct 2001.

[37] S. D. Pinches, H. L. Berk, D. N. Borba, B. N. Breizman, S. Briguglio, A. Fasoli, G. Fogaccia, M. P. Gryaznevich, V. Kiptily, M. J. Mantsinen, S. E. Sharapov, D. Testa, R. G. L. Vann, G. Vlad, F. Zonca, and JET-EFDA Contributors. The role of energetic particles in fusion plasmas. *Plasma Phys. Control. Fusion*, 46(12B):B187–B200, dec 2004.

[38] A. Ödblom, B. N. Breizman, S. E. Sharapov, T. C. Hender, and V. P. Pastukhov. Nonlinear magnetohydrodynamical effects in precessional fishbone oscillations. *Phys. Plasmas*, 9(1):155, jan 2002.

[39] W. W. Heidbrink. Beam-driven chirping instability in DIII-D. *Plasma Phys. Control. Fusion*, 37(9):937–949, sep 1995.

[40] V. N. Duarte, N. N. Gorelenkov, M. Schneller, E. D. Fredrickson, M. Podestà, and H. L. Berk. Study of the likelihood of Alfvénic mode bifurcation in NSTX and predictions for ITER baseline scenarios. *Nucl. Fusion*, 58(8):082013, aug 2018.

[41] J. B. Lestz, E. V. Belova, and N. N. Gorelenkov. Energetic-particle-modified global Alfvén eigenmodes. *Phys. Plasmas*, 25(4):042508, apr 2018.

[42] F. Wagner. A quarter-century of H-mode studies. *Plasma Phys. Control. Fusion*, 49(12B):B1–B33, dec 2007.

[43] M. D. Boyer, S. Kaye, and K. Erickson. Real-time capable modeling of neutral beam injection on NSTX-U using neural networks. *Nucl. Fusion*, 59(5):056008, may 2019.